

# Qualitative Concurrent Stochastic Games with Imperfect Information<sup>\*</sup>

Vincent Gripon and Olivier Serre

LIAFA (CNRS & Université Paris Diderot – Paris 7)

**Abstract.** We study a model of games that combines concurrency, imperfect information and stochastic aspects. Those are finite states games in which, at each round, the two players choose, *simultaneously* and *independently*, an action. Then a successor state is chosen accordingly to some fixed probability distribution depending on the previous state and on the pair of actions chosen by the players. Imperfect information is modeled as follows: both players have an equivalence relation over states and, instead of observing the exact state, they only know to which equivalence class it belongs. Therefore, if two partial plays are indistinguishable by some player, he should behave the same in both of them. We consider reachability (does the play eventually visit a final state?) and Büchi objective (does the play visit infinitely often a final state?). Our main contribution is to prove that the following problem is complete for 2-EXPTIME: decide whether the first player has a strategy that ensures her to almost-surely win against *any* possible strategy of her opponent. We also characterise those strategies needed by the first player to almost-surely win.

## 1 Introduction

Perfect information turn based two-player games on a graph [10] are widely studied in computer science. Indeed, they are a useful tool for both theoretical (for instance the modern proofs of Rabin’s complementation lemma rely on the memoryless determinacy of parity games [11]) and more practical applications. On the practical side, a major application of games is for the verification of reactive open systems. Those are systems composed of both a program and some (possibly hostile) environment. The verification problem consists of deciding whether the program can be restricted so that the system meets some given specification whatever the environment does. Here, restricting the system means synthesizing some controller, which, in term of games, is equivalent to designing a winning strategy for the player modeling the program [14].

The perfect information turn-based model, even if it suffices in many situations, is somewhat weak for the following two reasons. First, it does not permit to capture the behavior of real concurrent models where, in each step, the program and its environment *independently* choose moves, whose *parallel* execution

---

<sup>\*</sup> Supported by the ANR project JADE and by the ESF project GASICS.

determines the next state of the system. Second, in this model both players have, at each time, a perfect information on the current state of the play: this, for instance, forbids to model a system where the program and the environment share some public variables while having also their own private variables [15].

In this paper, we remove those two restrictions by considering concurrent stochastic games with imperfect information. Those are finite states games in which, at each round, the two players choose *simultaneously* and *independently* an action. Then a successor state is chosen accordingly to some fixed probability distribution depending on the previous state and on the pair of actions chosen by the players. Imperfect information is modeled as follows: both players have an equivalence relation over states and, instead of observing the exact state, they only see to which equivalence class it belongs. Therefore, if two partial plays are indistinguishable by some player, he should behave the same in both of them. Note that this model naturally captures several models studied in the literature [1, 9, 7, 8]. The winning conditions we consider here are reachability (is there a final state eventually visited?), Büchi (is there a final state that is visited infinitely often?) and their dual versions, safety and co-Büchi.

We study qualitative properties of those games (note that quantitative properties — *e.g.* deciding whether the value of the game is above a given threshold — are already undecidable in much weaker models [13]). More precisely, we investigate the question of deciding whether some player can almost-surely win, that is whether he has a strategy that wins with probability 1 against any counter strategy of the opponent. Our main contribution is to prove that, for both reachability and Büchi objectives, one can decide, in doubly exponential time (which is proved to be optimal), whether the first player has an almost-surely winning strategy. Moreover, when it is the case, we are also able to construct such a finite-memory strategy. We also provide intermediate new results concerning positive winning in safety (and co-Büchi)  $1\frac{1}{2}$ -player games (*a.k.a* partial observation Markov decision process).

**Related work.** Concurrent games with perfect information have been deeply investigated in the last decade [2, 1, 7]. Games with imperfect information have been considered for turn-based model [15] as well as for concurrent models with only one imperfectly informed player [9, 8]. To our knowledge, the present paper provides the first positive results on a model of games that combines concurrency, imperfect information (on both sides) and stochastic transition function. In a recent independent work [4], Bertrand, Genest and Gimbert obtain similar results than the one presented here for a closely related model. The main differences with our model are the following: Bertrand *et al.* consider a slightly weaker model of games in which the players may observe their own actions, and they allow the players to use richer strategies where the players can randomly update their memory (note that those strategies when used in our model seem strictly more powerful than the one we consider [12]). Bertrand *et al.* also discuss qualitative determinacy results and consider the case where a player is more informed than the other. We refer the reader to [4] for a detailed exposition.

## 2 Definitions

A **probability distribution** over a finite set  $X$  is a mapping  $d : X \rightarrow [0, 1]$  such that  $\sum_{x \in X} d(x) = 1$ . In the sequel we denote by  $\mathcal{D}(X)$  the set of probability distributions over  $X$ .

Given some set  $X$  and some equivalence relation  $\sim$  over  $X$ ,  $[x]_\sim$  stands for the equivalence class of  $x$  for  $\sim$  and  $X/\sim = \{[x]_\sim \mid x \in X\}$  denotes the set of equivalence classes of  $\sim$ .

For some finite alphabet  $A$ ,  $A^*$  (*resp.*  $A^\omega$ ) designates the set of finite (*resp.* infinite) words over  $A$ .

### 2.1 Arenas

A **concurrent arena with imperfect information** is a tuple  $\mathcal{A} = \langle S, \Sigma_E, \Sigma_A, \delta, \sim_E, \sim_A \rangle$  where

- $S$  is a finite set of control states;
- $\Sigma_E$  (*resp.*  $\Sigma_A$ ) is the (finite) set of actions for Eve (*resp.* Adam);
- $\delta : S \times \Sigma_E \times \Sigma_A \rightarrow \mathcal{D}(S)$  is the transition (total) function;
- $\sim_E$  and  $\sim_A$  are two equivalence relations over states.

A play in a such an arena proceeds as follows. First it starts in some initial state  $s$ . Then Eve picks an action  $\sigma_E \in \Sigma_E$  and, *simultaneously and independently*, Adam chooses an action  $\sigma_A \in \Sigma_A$ . Then a successor state is chosen accordingly to the probability distribution  $\delta(s, \sigma_E, \sigma_A)$ . Then the process restarts: the players choose a new pair of actions that induces, together with the current state, a new state and so on forever. Hence a **play** is an infinite sequence  $s_0 s_1 s_2 \dots$  in  $S^\omega$  such that for every  $i \geq 0$ , there exists  $(\sigma_E, \sigma_A) \in \Sigma_E \times \Sigma_A$  with  $\delta(s_i, \sigma_E, \sigma_A)(s_{i+1}) > 0$ . In the sequel we refer to a prefix of a play as a **partial play** and we denote by  $\text{Plays}(\mathcal{A})$  the set of all plays in arena  $\mathcal{A}$ .

The intuitive meaning of  $\sim_E$  (*resp.*  $\sim_A$ ) is that two states  $s_1$  and  $s_2$  such that  $s_1 \sim_E s_2$  (*resp.*  $s_1 \sim_A s_2$ ) cannot be distinguished by Eve (*resp.* by Adam). We easily extend the relation  $\sim_E$  to partial plays: let  $\lambda = s_0 s_1 \dots s_n$  and  $\lambda' = s'_0 s'_1 \dots s'_n$  be two partial plays, then  $\lambda \sim_E \lambda'$  if and only if  $s_i \sim_E s'_i$  for all  $i = 0, \dots, n$ .

Note that perfect information concurrent arenas (in the sense of [2, 1]) correspond to the special case where  $\sim_E$  and  $\sim_A$  are the equality relation over  $S$ .

### 2.2 Strategies

In order to choose their moves the players follow strategies, and, for this, they may use all the information they have about what was played so far. However, if two partial plays are equivalent for  $\sim_E$ , then Eve cannot distinguish them, and should therefore behave the same. This leads to the following notion.

An **observation-based strategy** for Eve is a function  $\varphi_E : (S/\sim_E)^* \rightarrow \mathcal{D}(\Sigma_E)$ , *i.e.*, to choose her next action, Eve considers the sequence of observations she got so far. In particular, a strategy  $\varphi_E$  is such that  $\varphi_E(\lambda) = \varphi_E(\lambda')$  whenever  $\lambda \sim_E \lambda'$ . Observation-based strategies for Adam are defined similarly.

Of special interest are those strategies that does not require memory: a **memoryless observation-based strategies** for Eve is a function from  $S/\sim_E \rightarrow \mathcal{D}(\Sigma_E)$ , that is to say these strategies only depend of the current equivalence class.

A **uniform strategy** for some player  $X$  is a strategy  $\varphi$  such that for all partial play  $\lambda$ , the probability measure  $\varphi(\lambda)$  is uniform, *i.e.*, for all action  $\sigma_X \in \Sigma_X$ , either  $\varphi(\lambda)(\sigma_X) = 0$  or  $\varphi(\lambda)(\sigma_X) = \frac{1}{|\{\sigma_X \in \Sigma_X | \varphi(\lambda)(\sigma_X) \neq 0\}|}$ . The set of memoryless uniform strategies for  $X$  is a finite set containing  $(2^{|\Sigma_X|} - 1)^{|\Sigma|}$  elements. Equivalently those strategies can be seen as functions to (non-empty) sets of (authorised) actions.

A **finite-memory strategy** for Eve with memory  $M$  ( $M$  being a finite set) is some triple  $\varphi = (Move, Up, m_0)$  where  $m_0 \in M$  is the initial memory,  $Move : M \rightarrow \mathcal{D}(\Sigma_E)$  associates a distribution of actions with any element in the memory  $M$  and  $Up : M \times S/\sim_E \rightarrow M$  is a mapping updating the memory with respect to some observation. One defines  $\varphi(s_0) = Move(m_0)$  and  $\varphi(s_0 \cdots s_n) = Move(Up(\cdots Up(Up(m_0, [s_1]/\sim_E), [s_2]/\sim_E), \cdots, [s_n]/\sim_E) \cdots)$  for any  $n \geq 1$ . Hence, a finite-memory strategy is some observation-based strategy that can be implemented by a finite transducer whose set of control states is  $M$ .

*Remark 1.* Note that in our definition of a strategy (and more generally in the definition of a play) we implicitly assume that the players only observe the sequence of states and not the corresponding sequence of actions. While the fact that Eve does not observe what Adam played is rather fair (otherwise imperfect information on states would make less sense) one could object that Eve should observe the actions she played so far. Here, our view of a (randomised) strategy is the following: when Eve respects some strategy, it means that whenever she has to play, her strategy provides her a distribution that she sends to some scheduler that, together with the distribution chosen by Adam, picks the next state. Indeed, it permits for instance to model a system in which some agent does not have the resources to implement himself randomisation.

An alternative option would be to consider that Eve flips a coin to pick her action and then sends this action to the scheduler that, together with the action chosen by Adam, picks the next state. In this case, a strategy should depend on the sequence of states together with the associated sequence of actions played by Eve. We argue that this second approach can be simulated easily by the first one, hence justifying our initial choice. Indeed, one can always enrich the set of states to encode the last pair of actions played and then use the equivalence relations  $\sim_E / \sim_A$  to hide / show part of this information to the respective players.

### 2.3 Probability Space and Outcomes of Strategies

Let  $\mathcal{A} = \langle S, \Sigma_E, \Sigma_A, \delta, \sim_E, \sim_A \rangle$  be a concurrent arena with imperfect information, let  $s_0 \in S$  be an initial state,  $\varphi_E$  be a strategy for Eve and  $\varphi_A$  be a strategy for Adam. In the sequel we are interested in defining the probability of a (measurable) set of plays knowing that Eve (*resp.* Adam) plays accordingly  $\varphi_E$  (*resp.*  $\varphi_A$ ). This is done in the classical way: first one defines the probability measure for basic sets of plays (called here *cones* and corresponding to plays having some initial common prefix) and then extends it in a unique way to all measurable sets.

First define  $Outcomes(s_0, \varphi_E, \varphi_A)$  to be the set of all possible plays when the game starts on  $s_0$  and when Eve and Adam plays respectively accordingly to  $\varphi_E$  and  $\varphi_A$ . More formally, an infinite play  $\lambda = s_0 s_1 \dots$  belongs to  $Outcomes(s_0, \varphi_E, \varphi_A)$  if and only if, for every  $i \geq 0$ , there is a pair of actions  $(\sigma_E, \sigma_A) \in \Sigma_E \times \Sigma_A$  with  $\delta(s_i, \sigma_E, \sigma_A)(s_{i+1}) > 0$  and s.t.  $\varphi_E(s_0 s_1 \dots s_i)(\sigma_E) > 0$  and  $\varphi_A(s_0 s_1 \dots s_i)(\sigma_A) > 0$  (*i.e.*  $\sigma_X$  is possible accordingly to  $\varphi_X$ , for  $X = E, A$ ).

Now, for any partial play  $\lambda$ , the **cone** for  $\lambda$  is the set  $cone(\lambda) = \lambda \cdot S^\omega$  of all infinite plays with prefix  $\lambda$ . Denote by  $Cones$  the set of all possible cones and let  $\mathcal{F}$  be the Borel  $\sigma$ -field generated by  $Cones$  considered as a set of basic open sets (*i.e.*  $\mathcal{F}$  is the smallest set containing  $Cones$  and closed under complementation, countable union and countable intersection). Then  $(Plays(\mathcal{A}), \mathcal{F})$  is a  $\sigma$ -algebra.

A pair of strategies  $(\varphi_E, \varphi_A)$  induces a probability space over  $(Plays(\mathcal{A}), \mathcal{F})$ . Indeed one can define a measure  $\mu_{s_0}^{\varphi_E, \varphi_A} : Cones \rightarrow [0, 1]$  on cones (this task is easy as a cone is uniquely defined by a finite partial play) and then uniquely extend it to a probability measure on  $\mathcal{F}$  using the Carathéodory Unique Extension Theorem. For this, one defines  $\mu_{s_0}^{\varphi_E, \varphi_A}$  inductively on cones:

- $\mu_{s_0}^{\varphi_E, \varphi_A}(s) = 1$  if  $s = s_0$  and  $\mu_{s_0}^{\varphi_E, \varphi_A}(s) = 0$  otherwise.
- For every partial play  $\lambda$  ending in some vertex  $s$ ,

$$\mu_{s_0}^{\varphi_E, \varphi_A}(\lambda \cdot s') = \mu_{s_0}^{\varphi_E, \varphi_A}(\lambda) \cdot \sum_{(\sigma_E, \sigma_A)} \varphi_E(\lambda)(\sigma_E) \cdot \varphi_A(\lambda)(\sigma_A) \cdot \delta(s, \sigma_E, \sigma_A)(s')$$

Denote by  $\Pr_{s_0}^{\varphi_E, \varphi_A}$  the unique extension of  $\mu_{s_0}^{\varphi_E, \varphi_A}$  to a probability measure on  $\mathcal{F}$ . Then  $(Plays(\mathcal{A}), \mathcal{F}, \Pr_{s_0}^{\varphi_E, \varphi_A})$  is a probability space.

### 2.4 Objectives, Value of a Game

Fix a concurrent arena with imperfect information  $\mathcal{A}$ . An objective for Eve is a measurable set  $\mathcal{O} \subseteq Plays(\mathcal{A})$ : a play is won by her if it belongs to  $\mathcal{O}$ ; otherwise it is won by Adam. A **concurrent game with imperfect information** is a triple  $(\mathcal{A}, s_0, \mathcal{O})$  where  $\mathcal{A}$  is a concurrent arena with imperfect information,  $s_0$  is an initial state and  $\mathcal{O}$  is an objective. In the sequel we focus on the following special classes of objectives (note that all of them are Borel sets hence measurable) that we define as means of a subset  $F \subseteq S$  of **final states**.

- A **reachability objective** is of the form  $S^*FS^\omega$ : a play is winning if it eventually goes through some final state.
- A **safety objective** is the dual of a reachability objective, *i.e.* is of the form  $(S \setminus F)^\omega$ : a play is winning if it never goes through a final state.
- A **Büchi objective** is of the form  $\bigcap_{k \geq 0} S^k S^*FS^\omega$ : a play is winning if it goes infinitely often through final states.
- A **co-Büchi objective** is the dual of a Büchi objective, *i.e.* is of the form  $S^*(S \setminus F)^\omega$ : a play is winning if it goes finitely often through final states.

A reachability (*resp.* safety, Büchi, co-Büchi) game is a game equipped with a reachability (*resp.* safety, Büchi, co-Büchi) objective. In the sequel we may replace  $\mathcal{O}$  by  $F$  when it is clear from the context which winning condition we consider.

Fix a concurrent game with imperfect information  $\mathbb{G} = (\mathcal{A}, s_0, \mathcal{O})$ . A strategy  $\varphi_E$  for Eve is **almost-surely winning** if, for any counter-strategy  $\varphi_A$  for Adam,  $\Pr_{s_0}^{\varphi_E, \varphi_A}(\mathcal{O}) = 1$ . If such a strategy exists, we say that Eve **almost-surely wins**  $\mathbb{G}$ . A strategy  $\varphi_E$  for Eve is **positively winning** if, for any counter-strategy  $\varphi_A$  for Adam,  $\Pr_{s_0}^{\varphi_E, \varphi_A}(\mathcal{O}) > 0$ . If such a strategy exists, we say that Eve **positively wins**  $\mathbb{G}$ .

### 3 Knowledge Arena

For the rest of this section we let  $\mathcal{A}$  be a concurrent arena with imperfect information with  $\mathcal{A} = \langle S, \Sigma_E, \Sigma_A, \delta, \sim_E, \sim_A \rangle$  and let  $s_0 \in S$  be some initial state.

Remark that in our model, the players do not observe the actions they play but they may know the distribution they have chosen. Therefore, one could consider a new arena in which the states have a component indicating the domain of the last distribution chosen by Eve, and that this component is visible only to her (it is hidden to Adam by the equivalence relation  $\sim_A$ ).

Even, if she does not see the precise control state, Eve can deduce information about it from previous information on the control state and from the set of possible actions she just played. We should refer to this as the **knowledge** of Eve, which formally is a set of states. Assume Eve knows that the current state belongs to some set  $K \subseteq S$ . After the next move Eve observes the equivalence class  $[s]_{\sim_E}$  of the new control state and she also knows the subset  $D$  of actions she may have played (it is the domain of the distribution she chose): hence she can compute the set of possible states the play can be in. This is done using the function  $\text{UpKnow} : 2^S \times [S]_{\sim_E} \times 2^{\Sigma_E} \rightarrow 2^{\Sigma_S}$  defined by letting

$$\begin{aligned} \text{UpKnow}(K, [s]_{\sim_E}, D) = \\ \{t \sim_E s \mid \exists r \in K, \sigma'_E \in D, \sigma_A \in \Sigma_A \text{ s.t. } \delta(r, \sigma'_E, \sigma_A)(t) > 0\}, \end{aligned}$$

*i.e.* in order to update her current knowledge, observing in which equivalence class is the new control state, and knowing that she played an action in  $D \subseteq \Sigma_E$ , Eve computes the set of all states in this class that may be reached from a state in her former knowledge.

Based on our initial remark and on the notion of knowledge we define the **knowledge arena associated with  $\mathcal{A}$** , denoted  $\mathcal{A}^K$ . The arena  $\mathcal{A}^K$  is designed to make explicit the information Eve can collect on her moves (*i.e.* the domain of the distributions she plays) and on the possible current state. We define  $\mathcal{A}^K = \langle S^K, \Sigma'_E, \Sigma_A, \delta^K, \sim_E^K, \sim_A^K \rangle$  as follows:

- $S^K = \{(s, K, D) \in S \times 2^S \times 2^{\Sigma_E} \mid K \subseteq [s]_{\sim_E}\}$ : the first component is the real state, the second one is the current knowledge of Eve and the third one is the domain of the last distribution she played;
- $\Sigma'_E = \Sigma_E \times (2^{\Sigma_E} \setminus \emptyset)$ : actions of Eve will now contain information on the domain of actions of the distributions she picks;
- $\delta^K((s, K, D), (\sigma_E, D'), \sigma_A)(s', K', D'') = 0$  if  $D' \neq D''$  or  $K' \neq \text{UpKnow}(K, [s']_{\sim_E}, D')$ ; and  $\delta^K((s, K, D), (\sigma_E, D'), \sigma_A)(s', K', D') = \delta(s, \sigma_E, \sigma_A)(s')$  otherwise:  $\delta^K$  behaves as  $\delta$  on the first components and *deterministically* updates both the knowledge and the information on the domain;
- $(s, K, D) \sim_E^K (s', K', D')$  if and only if  $K = K'$  (implying  $s \sim_E s'$ ) and  $D = D'$ : Eve observes her knowledge and the domain of her last distribution;
- $(s, K, D) \sim_A^K (s', K', D')$  if and only if  $s \sim_A s'$ .

The intuitive meaning of the enriched alphabet  $\Sigma'_E$  of Eve is that instead of choosing a distribution  $d : \Sigma_E \rightarrow [0, 1]$  Eve makes the domain  $\text{Dom} = \{s \in S \mid d(s) > 0\}$  of  $d$  explicit by choosing the distribution  $d^K$  where  $d^K(s, D) = d(s)$  if  $D = \text{Dom}$  and  $d^K(s, D) = 0$  otherwise. We call such a distribution  $d^K$  **well-formed** (*i.e.*  $d^K$  is obtained from some distribution  $d$  as just explained) and, in the sequel, whenever referring to strategies of Eve in  $\mathcal{A}^K$ , we will mean functions from sequences of observations into *well-formed* distributions.

Consider an observation-based strategy  $\varphi$  for Eve in the arena  $\mathcal{A}$ . Then it can be converted into an observation-based strategy on the associated knowledge arena. For this, remark that in the knowledge arena, those states reachable from the initial state  $(s_0, \{s_0\}, \emptyset)$  are of the form  $(s, K, D)$  with all states in  $K$  being equivalent with  $s$  with respect to  $\sim_E$ . Then one can define  $\varphi^K((s_0, K_0, D_0)(s_1, K_1, D_1) \cdots (s_n, K_n, D_n))$  as  $d^K$  where  $d = \varphi([s_0]_{\sim_E}[s_1]_{\sim_E} \cdots [s_n]_{\sim_E})$  is the corresponding distribution given by  $\varphi$ . Note that  $\varphi^K$  is observation-based as, for all  $0 \leq h \leq n$ ,  $[s_h]_{\sim_E}$  is uniquely defined from the  $K_h$ , that are observed by Eve in the knowledge arena.

Conversely, any observation-based strategy in the knowledge arena can be converted into an observation-based strategy in the original arena. Indeed, consider some observation-based strategy  $\varphi^K$  in the knowledge arena: it is a mapping from  $(2^S \times 2^{\Sigma})^*$  into  $\mathcal{D}(\Sigma'_E)$  (the equivalent classes of the relation  $\sim_E^K$  are, by definition, isomorphic with  $2^S \times 2^{\Sigma}$ ). Now, note that Eve can, while playing in  $\mathcal{A}$ , remember the domain of the distributions she played and compute on the fly her current knowledge (applying function  $\text{UpKnow}$  to her previous knowledge and to the domain of the last distribution played): hence along a play  $s_0 s_1 \cdots s_n$  she can compute the corresponding sequence  $(K_0, D_0)(K_1, D_1) \cdots (K_n, D_n)$  of knowledge / domain. Now it suffices to consider the observation-based strategy  $\varphi$  for Eve in the initial arena defined by:

$$\varphi(s_0 s_1 \cdots s_n) = \varphi^K((K_0, D_0)(K_1, D_1) \cdots (K_n, D_n))$$

Note that this last transformation (taking a strategy  $\varphi^K$  and producing a strategy  $\varphi$ ) is the inverse of the first transformation (taking a strategy  $\varphi$  and producing a strategy  $\varphi^K$ ). In particular, it proves that the observation-based strategies in both arena are in bijection. It should be clear that those strategies for Adam in both games are the same (as what he observes is identical).

Assume that  $\mathcal{A}$  is equipped with a set  $F$  of final states. Then one defines the final states in  $\mathcal{A}^K$  by letting  $F^K = \{(f, K, D) \mid f \in F\} \cap S^K$ : this allows to define an objective  $\mathcal{O}^K$  in  $\mathcal{A}^K$  from an objective  $\mathcal{O}$  in  $\mathcal{A}$ . Based on the previous observations, we derive the following.

**Proposition 1.** *Let  $\mathbb{G} = (\mathcal{A}, s_0, \mathcal{O})$  be some imperfect information game equipped with a reachability (resp. safety, Büchi, co-Büchi) objective. Let  $\mathbb{G}^K = (\mathcal{A}^K, (s_0, \{s_0\}, \emptyset), \mathcal{O}^K)$  be the associated game played on the knowledge arena. Then for any strategies  $\varphi_E, \varphi_A$  for Eve and Adam, the following holds:*

*$\Pr_{s_0}^{\varphi_E, \varphi_A}(\mathcal{O}) = \Pr_{(s_0, \{s_0\}, \emptyset)}^{\varphi_E^K, \varphi_A^K}(\mathcal{O}^K)$ . In particular, Eve has an almost-surely winning observation-based strategy in  $\mathbb{G}$  if and only if she has one in  $\mathbb{G}^K$ .*

In the setting of the previous proposition, consider the special case where Eve has an almost-surely winning observation-based strategy  $\varphi^K$  in  $\mathbb{G}^K$  that only depends on the current knowledge (in particular, it is *memoryless*). Then the corresponding almost-surely winning observation-based strategy  $\varphi$  in  $\mathbb{G}$  is, in general, not memoryless, but can be implemented by a finite transducer whose set of control states is precisely the set of possible knowledges for Eve. More precisely the strategy consists in computing and updating on the fly (using a finite automaton) the value of the knowledge after the current partial play and to pick the next action by solely considering the knowledge. We may refer at such a strategy  $\varphi$  as a **knowledge-only strategy**.

## 4 Decidability Results

### 4.1 Reachability Objectives

The main result of this section is the following.

**Theorem 1.** *For any reachability concurrent game with imperfect information, one can decide, in doubly exponential time, whether Eve has an almost-surely winning strategy. If Eve has such a strategy then she has a knowledge-only uniform strategy, and such a strategy can be effectively constructed.*

Before proving Theorem 1 we first establish an intermediate result. A concurrent game (with imperfect information) in which one player has only a single available action is what we refer as a  **$1\frac{1}{2}$ -player game with imperfect information** (those games are also known in the literature as *partially observable Markov Decision Processes*). The following result is a key ingredient for the proofs of Proposition 2 and Theorem 1.



**Lemma 1.** *Consider an  $1\frac{1}{2}$ -player safety game with imperfect information. Assume that the player has an observation-based strategy that is positively winning. Then she also has an observation-based finite memory strategy that is positively winning. Moreover, both the strategy and the set of positively winning states can be computed in time  $\mathcal{O}(2^{|S|})$ .*

*Proof (Sketch).* Consider the knowledge arena and call a knowledge  $K$  *surely winning* if the player has a knowledge based strategy that is surely winning from any  $(s, K, D)$  with  $s \in K$  and  $D \subseteq \Sigma_E$ . We prove, that if the player has a positively winning strategy, then the set of winning knowledges is non empty and that it comes with a memoryless surely winning strategy (that consists in staying in the surely winning component). This set also contains at least a singleton  $\{s\}$  (meaning that if the player knows that she is in  $s$  then she can surely win): call such states  $s$  *surely winning*. Then, one proves that positively winning states are exactly those that are connected (in the graph sense) to some surely winning state by a path made of *non-final* states. Hence a positively winning strategy consists in playing some initial actions randomly (trying to reach a surely winning state) and then in mimicking a knowledge-only surely winning strategy. Complexity comes with a fixpoint definition of the previous objects.  $\square$

Fix, for the rest of this section, a concurrent game with imperfect information  $\mathbb{G} = (\mathcal{A}, s_0, \mathcal{O})$  equipped with a reachability objective  $\mathcal{O}$  defined from a set  $F$  of final states. We set  $\mathcal{A} = \langle S, \Sigma_E, \Sigma_A, \delta, \sim_E, \sim_A \rangle$ . We also consider  $\mathbb{G}^K = (\mathcal{A}^K, (s_0, \{s_0\}, \emptyset), \mathcal{O}^K)$  to be the corresponding knowledge game.

To prove Theorem 1, one first defines (in a non constructive way) a knowledge-only uniform strategy  $\varphi$  for Eve as follows. We let

$$\mathcal{K}^{\text{AS}} = \{K \in 2^S \mid \exists \varphi_E \text{ knowledge-based strategy for Eve s.t. } \varphi_E \text{ is almost-surely winning for Eve in } \mathbb{G}^K \text{ from any } (s, K, D) \text{ with } s \in K \text{ and } D \subseteq \Sigma_E\}$$

be the set of knowledges made only by almost-surely winning states for Eve (note here that we require that the almost-surely winning strategy is the same for all configurations with the same knowledge).

One can prove that, from a configuration with knowledge  $K \in \mathcal{K}^{\text{AS}}$ , Eve always has at least one action which ensures that she remains in  $\mathcal{K}^{\text{AS}}$ , and we define  $\varphi$  as the knowledge-only uniform strategy that chooses at random one of these safe actions. The next proposition shows that  $\varphi$  is almost-surely winning for Eve.

**Proposition 2.** *The strategy  $\varphi$  is almost-surely winning for Eve from states whose Eve's knowledge is in  $\mathcal{K}^{\text{AS}}$ .*

*Proof (sketch).* To prove that  $\varphi$  is almost-surely winning, one needs to prove that it is almost surely-winning against any strategy of Adam. However, once  $\varphi$  is fixed (and as it is a knowledge-only strategy), one gets  $1\frac{1}{2}$ -player game in which only Adam is making choices. Proving that  $\varphi$  is almost surely winning is

therefore equivalent to proving that Adam cannot positively win in this new game (for a safety objective). For this we use Lemma 1 to argue that it suffices to prove that  $\varphi$  is winning against any finite-memory strategy of Adam. This fact permits us to conclude.  $\square$

Now one can prove Theorem 1. First Eve almost-surely wins in  $\mathbb{G}$  if and only if she almost-surely wins in  $\mathbb{G}^K$  if and only if  $\{s_0\} \in \mathcal{K}^{\text{AS}}$ , *i.e.* (using Proposition 2) if and only if Eve has a knowledge-only uniform strategy in  $\mathbb{G}^K$ . Now, to decide whether Eve almost-surely wins  $\mathbb{G}$ , it suffices to check, for any possible knowledge-only uniform strategy  $\varphi$  for her, whether it is almost-surely winning. Once  $\varphi$  is fixed, it leads, from Adam's point of view, to a  $1\frac{1}{2}$ -player safety game  $\mathbb{G}_\varphi$  where the player positively wins if and only if  $\varphi$  is not almost-surely winning. Hence Lemma 1 implies that deciding whether  $\varphi$  is almost-surely winning can be done in time exponential in the size of  $\mathbb{G}_\varphi$ , which itself is of exponential size in  $|S|$ . Hence deciding whether a knowledge-only uniform strategy for Eve is winning can be done in doubly exponential time (in the size of  $|S|$ ). The set of knowledge-only uniform strategies for Eve is finite and its size is doubly exponential in the size of the game. Hence the overall procedure, that tests every possible such strategies, requires doubly exponential time. As effectivity is immediate, this concludes the proof of Theorem 1.

The naive underlying algorithm of Theorem 1 turns out to be optimal.

**Theorem 2.** *Deciding whether Eve almost-surely wins a concurrent game with imperfect information is a 2-EXPTIME-complete problem.*

*Proof (sketch).* The proof is a generalisation of a similar result given in [8] showing EXPTIME-hardness of concurrent games *only one* player is imperfectly informed. The idea is to simulate an alternating exponential space Turing machine (without input). We design a game where the players describe the run of such a machine: transitions from existential (*resp.* universal) states are chosen by Eve (*resp.* Adam) and Adam is also in charge of describing the successive configurations of the machine. To prevent him from cheating, Eve can secretly mark a cell of the tape, and latter check whether it was correctly updated (if not she wins). As she cannot store the exact index of the cell (it is of exponential size), she could cheat in the previous phase: hence Adam secretly marks some bit and one recall the value of the corresponding bit of the index of the marked cell: this bit is checked when Eve claims that Adam cheated (if it is wrong then she is loosing). Eve also wins if the described run is accepting. Eve can also restart the computation whenever she wants (this is useful when she cannot prove that Adam cheated): hence if the machine accepts the only option for Adam is to cheat, and Eve will eventually catch him with probability one. Now if the machine does not accept, the only option for Eve is to cheat, but it will be detected with positive probability.  $\square$

## 4.2 Büchi Objectives

We now consider the problem of deciding whether Eve almost-surely wins a Büchi game. The results and techniques are similar to the one for reachability

games. In particular, we need to establish the following intermediate result (the proof is very similar to the one of Lemma 1 except that now the winning states are those connected by *any kind* of path to a surely winning state).

**Lemma 2.** *Consider an  $1\frac{1}{2}$ -player co-Büchi game with imperfect information. Assume that the player has an observation-based strategy that is positively winning. Then she also has an observation-based finite memory strategy that is positively winning. Moreover, both the strategy and the set of positively winning states can be computed in time  $\mathcal{O}(2^{|S|})$ .*

From Lemma 2 and extra intermediate results we derive our main result. Again, the key idea is to prove that the strategy that plays randomly inside the almost-surely winning region is an almost-surely winning strategy.

**Theorem 3.** *For any Büchi concurrent game with imperfect information, one can decide, in doubly exponential time, whether Eve has an almost-surely winning strategy. If Eve has such a strategy then she has a knowledge-based uniform memoryless strategy, and such a strategy can be effectively constructed. The doubly exponential time complexity bound is optimal.*

## 5 Discussion

The main contribution of this paper is to prove that one can decide whether Eve has an almost-surely winning strategy in a concurrent game with imperfect information equipped with a reachability objective or a Büchi objective.

A natural question is whether this result holds for other objectives, in particular for co-Büchi objectives. In a recent work [3], Baier *et al.* established undecidability of the emptiness problem for probabilistic Büchi automata on infinite words. Such an automaton can be simulated by a  $1\frac{1}{2}$ -player imperfect information game: the states of the game are the one of the automaton, they are all equivalent for the player, and therefore an observation based strategy is an infinite word. Hence a pure (*i.e.* non-randomised) strategy in such a game coincide with an input word for the automaton. From this fact, Baier *et al.* derived that it is undecidable whether, in a  $1\frac{1}{2}$ -player co-Büchi game with imperfect information, Eve has an almost-surely winning *pure* strategy.

One can also consider the stochastic-free version of this problem (an arena is **deterministic** iff  $\delta(q, \sigma_E, \sigma_A)(q') \in \{0, 1\}$  for all  $q, q', \sigma_E, \sigma_A$ ) and investigate whether one can decide if Eve has an almost-surely winning strategy in a *deterministic* game equipped with a co-Büchi objective. We believe that the  $1\frac{1}{2}$ -player setting can be reduced to this new one, hence allowing to transfer undecidability results [12]. An even weaker model to consider is the stochastic-free model in which Adam has perfect information about the play [8].

It may happen that Eve has no almost-surely winning strategy while having a family  $(\varphi_\varepsilon)_{0 < \varepsilon < 1}$  of strategies such that  $\varphi_\varepsilon$  ensures to win with probability at least  $1 - \varepsilon$ . Such a family is called **limit-surely winning**. Deciding existence of such families is a very challenging problem: indeed, in many practical situations,

it is satisfying enough if one can control the risk of failing. Even if those questions have been solved for perfect information games [1], as far as we know, there has not been yet any result obtained in the imperfect information setting.

Even if the algorithms provided in this paper are "optimal", they are rather naive (checking all strategies for Eve may cost a lot in practice). Hence, one should look for fixpoint-based algorithms as the one studied in [8]: it would be of great help for a symbolic implementation, and it could also be a useful step toward a solution of the problem of finding limit-surely winning strategies. Note that there are already efficient techniques and tools for finding *sure* winning strategies in subclasses of concurrent games with imperfect information [6, 5].

**Acknowledgements.** The authors want to thank the anonymous reviewers as well as Florian Horn for their useful comments on preliminary versions of this work.

## References

1. L. de Alfaro and T.A. Henzinger. Concurrent omega-regular games. In *Proceedings of LICS'00*, pages 141–154, 2000.
2. L. de Alfaro, T.A. Henzinger, and O. Kupferman. Concurrent reachability games. *Theoretical Computer Science*, 386(3):188–217, 2007.
3. C. Baier, N. Bertrand, and M. Größer. On decision problems for probabilistic büchi automata. In *Proceedings of FoSSaCS 2008*, volume 4962 of *LNCS*, pages 287–301. Springer, 2008.
4. N. Bertrand, B. Genest, and H. Gimbert. Qualitative Determinacy and Decidability of Stochastic Games with Signals. In *Proceedings of LICS 2009*. To appear.
5. D. Berwanger, K. Chatterjee, M. De Wulf, L. Doyen, and T.A. Henzinger. Alpaga: A tool for solving parity games with imperfect information. In *Proceedings of TACAS 2009*, volume 5505 of *LNCS*, pages 58–61. Springer, 2009.
6. D. Berwanger, K. Chatterjee, L. Doyen, T.A. Henzinger, and S. Raje. Strategy construction for parity games with imperfect information. In *Proceedings of CONCUR 2008*, volume 5201 of *LNCS*, pages 325–339. Springer, 2008.
7. K. Chatterjee. *Stochastic  $\omega$ -Regular Games*. PhD thesis, University of California, 2007.
8. K. Chatterjee, L. Doyen, T.A. Henzinger, and J.-F. Raskin. Algorithms for omega-regular games with imperfect information. *Logical Methods in Computer Science*, 3(3), 2007.
9. K. Chatterjee and T.A. Henzinger. Semiperfect-information games. In *Proceedings of FST&TCS 2005*, volume 3821 of *LNCS*, pages 1–18. Springer, 2005.
10. E. Grädel, W. Thomas, and Th. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *LNCS*. Springer, 2002.
11. Y. Gurevich and L. Harrington. Trees, automata, and games. In *Proceedings of STOC 1982*, pages 60–65, 1982.
12. F. Horn. Private communication. February 2009.
13. A. Paz. *Introduction to probabilistic automata*. Academic Press New York, 1971.
14. P.J. Ramadge and W.M. Wonham. Supervisory Control of a Class of Discrete Event Processes. *SIAM Journal on Control and Optimization*, 25:206, 1987.
15. J.H. Reif. The complexity of two-player games of incomplete information. *Journal of Computer and System Sciences*, 29(2):274–301, 1984.